Evaluating the Impact of Fault Recovery on Superscalar Processor Performance

Sato, Toshinori
System LSI Research Center, Kyushu University

Chiyonobu, Akihiro
Graduate School of Computer Science and System Engineering, Kyushu Institute of Technology

<https://doi.org/10.15017/6353>

出版情報 : Proc. of 12th Pacific Rim International Symposium on Dependable Computing, pp.369-370, 2006-12. Pacific Rim International Symposium on Dependable Computing
バージョン :
権利関係 :

Evaluating the Impact of Fault Recovery on Superscalar Processor Performance⁺

Toshinori Sato
System LSI Research Center
Kyushu University
toshinori.sato@computer.org

Akihiro Chiyonobu
Graduate School of Computer Science and System Engineering
Kyushu Institute of Technology
chiyo@mickey.ai.kyutech.ac.jp

1. Introduction

Current semiconductor technologies have become susceptible to high-energy neutrons from space. Following the trends in smaller transistors, lower supply voltage, and higher clock frequency, current microprocessors are susceptible to soft errors, which constitute the vast majority of hardware failures. Based on these trends, it is expected that the quality with respect to reliability becomes important as well as performance for microprocessors.

In light of this, a lot of fault-tolerance microarchitectures are recently proposed (see Section 2 in [1]). These studies mainly focus on detecting transient faults, and hence almost every previous study evaluated processor performance in the absence of faults. This analysis only presents the performance impact of constraints introduced by fault detection mechanism. While several papers propose fault recovery mechanisms, only a few evaluated their impact on performance [2, 3]. One of the reasons why this evaluation methodology is widely selected is that faults are expected to be rare enough that the overall performance will be determined by fault-free behavior. However, evaluating recovery cost of fault tolerant execution is also important, because it is predicted that transient hardware faults occur more frequently as semiconductor technology is improved. Therefore, this paper focuses on recovery from faults.

2. Fault Tolerance Mechanism

We are investigating the use of instruction reissue technique as a hardware mechanism to detect and recover from transient faults [4]. Originally, the instruction reissue mechanism is proposed for incorrect data speculation recovery. We modified and applied this mechanism for fault-tolerance. Since future microprocessors will utilize data speculation and include the instruction reissue mechanism, this fault-tolerance mechanism costs the least hardware overhead. We have already proposed two transparent fault recovery mechanisms [4], each of which uses recovery mechanisms for speculative execution. One of them uses the recovery mechanism for mispredicted branches, and the other uses that for misspeculated data dependences. In this paper, we will evaluate the performance impact introduced by the recovery mechanisms and compare them.

3. Simulation Results

We evaluate 4- and 8-way OoO execution superscalar processors. The SPEC2000 benchmark suite is used for this study. The details of our evaluation environment can be found in [1]. We use a 1K-entry reuse buffer in order to mitigate performance penalty due to redundant execution.

Table 1. %Increase in execution cycles without recovery

program	4-way	8-way
164.gzip	33.3	17.8
175.vpr	32.6	14.7
176.gcc	23.4	11.0
186.crafty	37.5	23.1
197.parser	28.1	13.5
252.eon	40.9	33.0
255.vortex	37.1	23.3
256.bzip2	53.7	45.7

Table 1 presents the percentage increase in execution cycles when the proposed fault-tolerance mechanism is introduced. Here, recovery from faults is not considered.

The fault injection module in our simulator randomly corrupts some instructions. Figure 1 shows the percentage increase in execution cycles when faults are injected. The horizontal line indicates the average fault frequency per one million cycles, while the vertical line indicates the corresponding percentage increase. There are four graphs in the figures. *Flush-* and *Selective-* mean that their corresponding fault recovery mechanisms are based on the recovery mechanism for control and data misspeculation, respectively. *-4* and *-8* mean that their corresponding processor models are 4- and 8-way superscalar processors, respectively. For both processor models, no significant difference between *Flush-* and *Selective-* mechanisms is found until the number of faults reaches about 1000 per million cycles. While the performance loss is increased over the point where approximately 1000 faults occur per million cycles, the fault frequency is much higher than the actual fault frequencies we intend.

4. Conclusions

This is the first paper comparing *Flush-* and *Selective-* recovery mechanisms for soft error tolerance. The detailed simulations do not show any considerable differences in recovery cost between two mechanisms. In addition, across the fault frequencies we intend, it is possible to ignore the cost, because the overall processor performance is dominated by fault-free behavior.

References

- [1] T. Sato: Exploiting instruction redundancy for transient fault tolerance, IEEE 18th DFT (2003)
- [2] T. M. Austin: DIVA: a reliable substrate for deep submicron microarchitecture design, IEEE/ACM 32nd MICRO (1999)
- [3] J. Ray, et al.: Dual use of superscalar datapath for transient-fault detection and recovery, IEEE/ACM 34th MICRO (2001)
- [4] T. Sato: A transparent transient faults tolerance mechanism for superscalar processors, IEICE Trans. Inf. & Sys., **E86-D(12)** (2003)

⁺ This work is partially supported by Grants-in-Aid for Scientific Research #16300019 and #176549 from Japan Society for the Promotion of Science.

Figure 1. Impact of the number of faults per million cycles on the percent increase in cycles